

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
11 April 2002 (11.04.2002)

PCT

(10) International Publication Number
WO 02/30061 A1

(51) International Patent Classification⁷: **H04L 12/56,**
H04Q 11/04

(21) International Application Number: **PCT/GB00/03783**

(22) International Filing Date: **3 October 2000 (03.10.2000)**

(25) Filing Language: **English**

(26) Publication Language: **English**

(71) Applicant (for all designated States except US): **DE-GREE2 INNOVATIONS LIMITED** [GB/GB]; 99 Charterhouse Street, London EC1M 6NQ (GB).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **DAVIES, Neil, James** [GB/GB]; 19 Chesnut Road, Long Ashton, Bristol BS41 9HR (GB). **THOMPSON, Peter, William**



[GB/GB]; 5 Normanton Road, Clifton, Bristol BS8 2TY (GB). **HOLYER, Judith, Yvonne** [GB/GB]; 55 Rowsham Mead, Hotwells, Bristol BS8 4YB (GB). **LAFAVE, Laura, Anne** [US/GB]; 10 Southernhay Crescent, Bristol BS8 4TT (GB). **VOWDEN, Christopher, James** [GB/GB]; 3 Stockwood Mews, St Annes Park, Bristol BS4 4QA (GB). **WILLMOTT, Graham** [GB/GB]; 78 Lyppiat Road, Redfield, Bristol BS5 9HS (GB).

(74) Agents: **WILLIAMS, David, John et al.**; Page White & Farrer, 54 Doughty Street, London WC1N 2LS (GB).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

[Continued on next page]

(54) Title: **FILTERING DATA FLOWS**



(57) Abstract: There is disclosed a method and apparatus for controlling an information flow in a data transmission system. A plurality of data packets are received. The plurality of data packets are selectively discarded and/or selectively time shifted in a predetermined manner. A modified output flow is thereby generated. The data packets of the modified output flow are selectively discarded and/or selectively time shifted in at least one of a plurality of further parallel steps. At least one further modified output flow is thereby generated. The modified output flow may be further modified in all of the plurality of further parallel steps. The modified output flow may be further modified in the selected one of the plurality of further parallel steps. The modified output flow may be further modified in selected ones of the plurality of further parallel steps.



(84) **Designated States (regional):** ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

5

FILTERING DATA FLOWSCross-Reference to Related Applications

This application is related to patent application number [Agent Refs: 200033, 200034, 200035, 200036, 200037, 200039] filed simultaneously hereto on 3
10 October 2000.

Field of the Invention

The present invention relates generally to the transmission of digital packets in
15 telecommunications systems, and particularly but not exclusively to contention management within and between streams of packets in store and forward networks.

Background of the Invention

20 There exist many techniques for controlling digital packet information flows. Some such techniques involve providing a certain quality of service for different types of traffic. There are a number of general requirements associated with ensuring quality of service levels. Where quality differentiation is provided for more than one stream, it is necessary to ensure that different quality requirements for the
25 various streams of packets are individually and collectively met within the constraint of the finite quality available. In order to provide differential levels of loss and delay as well as throughput, the quantity of input data and its temporal pattern need to be constrained. The quantity of input data serviced, i.e. the throughput of data, is known as long term management. Controlling the temporal
30 pattern of data is known as short term management. The traffic in particular streams also needs to be protected from the consequential effects of burstiness in other streams. For example, individual streams of traffic need to be protected from the effects of protocols such as TCP, which is designed to use as much bandwidth as possible without regard to other streams, and from malicious
35 intentions or errors in the end devices of the network. It is also necessary to manage the interleaving of individual streams within the constraints of the available resources.

5

The management of quality of service is particularly difficult at the periphery or edge of a network, as devices here are not under the control of the central administrator of the network. The behaviour of these devices therefore cannot be assumed or predicted. The ever-increasing diversity in applications, traffic, and convergence also complicates the management of quality of service. Different traffic has different quality of service requirements, and the consequences of delay and/or loss differs for different traffic according to the interpretation placed on them by an application.

10

15 The possibility of replicating network devices to enable traffic with different quality requirements to be physically separated and processed separately is impractical, as the implementation of network devices is expensive. For this reason, it is desirable to manage quality of service of different traffic using a single network device.

20

There is currently a trend towards forming traffic patterns as constant bit rate patterns, with the aim of increasing predictability. This is deterministic control which focuses on improving the loss characteristics and efficiency of a network.

However such techniques have the disadvantage that quality assurance in the presence of "over-booking" requires total global knowledge of the behaviour of the sources in the network such as their on/off times, their relative phase, and the effects of their past history on the state of the network.

25

Recent work by the Applicant has shown how to control and manage individual streams in their use of resources by themselves and in the competition for use of resources with others. This approach generates a class of useful mechanisms.

30

It is therefore an object of the present invention to provide an improved technique for controlling information flow in a data transmission system, which enables the control of the quality of service requirements for different types of traffic to be improved.

35

5 Summary of the Invention

According to the present invention there is provided a method of controlling an information flow in a data transmission system comprising; receiving a plurality of data packets; selectively discarding and/or selectively time shifting the plurality of data packets in a predetermined manner to thereby generate a modified output flow; selectively discarding and/or selectively time shifting the data packets of the modified output flow in at least one of a plurality of further parallel steps to thereby generate at least one further modified output flow.

Thus there is provided a generalization where composite shape streams can be combined together in an arbitrary fashion. Such composite mechanisms permit the management of properties of aggregate streams to limit, for example, a number of streams to an individual load bound and their combined load to another bound which may be lower than the sum of the individual bounds. The invention covers all feasible acyclic combinations of individual quality management devices and collective quality management devices.

The modified output flow may be further modified in all of the plurality of further parallel steps.

The modified output flow may be further modified in a selected one of the plurality of further parallel steps.

The filtered output flow may be further modified in selected ones of the plurality of further parallel steps.

The step of selectively discarding and/or selectively time shifting packets may further comprise generating a modified output flow in which the packets are variably spaced.

The present invention further provides a controller for controlling an information

5 flow in a data transmission system, the controller comprising: first control means
for selectively discarding and/or selectively time shifting a plurality of data packets
in a predetermined manner to thereby generate a modified information flow; and a
plurality of parallel second control means for selectively discarding and/or
selectively time shifting the modified information flow in a predetermined manner,
10 thereby to provide a further modified information flow.

At least one control means may include modifying means for generating the
modified information flow in which the packets are variably spaced. The modified
means may comprise a modulator. At least one control means may comprise a
15 policer/shaper.

Brief Description of the Figures

The invention will be best understood by way of example with reference to the
20 accompanying Figures, in which:

Figure 1 illustrates a block diagram of the architecture of a multiplexor in
accordance with an exemplary embodiment of the invention;

Figure 2 illustrates the principle of queuing in accordance with an exemplary
25 embodiment of the invention;

Figure 3 illustrates the principle of allocating priority levels in accordance with an
exemplary embodiment of the invention;

Figure 4 illustrates the principle of stochastically servicing packets in accordance
with an exemplary embodiment of the invention;

30 Figure 5 illustrates the principle of conceptual packet length in accordance with an
exemplary embodiment of the invention.

Figure 6(a) and 6(b) illustrate the principle of cascading shaper/policers in
accordance with an exemplary embodiment of the invention;

Figure 7 illustrates the principle of defining a transported load in dependence on
35 an offered load in an exemplary embodiment of the present invention;

- 5 Figure 8 illustrates an exemplary embodiment of a Policer/Shaper for implementing the present invention; Figure 9 illustrates an exemplary embodiment of a Cherish/Urgency Multiplexer for implementing the present invention; Figure 10 illustrates the load profile for an exemplary Policer/Shaper implementation.

10

Description of the Preferred Embodiment

- Referring to Figure 1, there is shown a diagram illustrating the architecture of a basic multiplexor in respect of which the operation of the present invention is described herein. It should be noted that the application of the present invention is not limited to the specific architecture shown in Figure 1, and from reading the following description the skilled person will appreciate the general applicability of the present invention. The exemplary multiplexor 100 of Figure 1 multiplexes streams of data in a packet-based network as will be described in further detail hereinbelow.

- The exemplary multiplexor generally designated by reference numeral 100, comprises an input interface 100, a stream identifier 102, a plurality of policer/shapers 104, a cherish/urgency multiplexer 106, a rate limiter 108, an output interface 110, a queue memory 112, and a queue memory manager 114.

- The multiplexor 100 receives as an input on lines 116 an information flow. The information flow comprises packets which may be associated with various different types of data traffic. The packets may be received in parallel or in series. The input on line 116 may receive parallel information flows. The input data stream 116 may contain packets associated with various different types of traffic. For the purposes of the described example, it is assumed the data stream includes a voice over IP stream, a block data transfer stream, and a telnet traffic stream.

35

5 The input interface 100 provides the basic functionality necessary to receive transmitted data packets from an external system or device via a transmission medium. The input interface 100 receives packets from the input data stream 116 and forwards the received packets, preferably in parallel, on lines 118 to both the stream identifier 102 and the queue memory manager 114. The input interface
10 must have functionality which is appropriate for the particular external system or device from which packets on the input data stream 116 originate. The input interface is not directly responsible for any processing related to the quality management of the data stream in accordance with the present invention. The input interface may have some basic functionality to perform an initial check on
15 packets received in the input data stream. For example, if a corrupted packet is received the input interface may discard the packet.

The structure of the input interface is implementation dependent. The basic functionality of the input interface, required to perform any desired features, will be
20 apparent to one skilled in the art.

The queue memory manager 114 is responsible for managing a set of queues and a packet storage area. The queue memory manager 114 receives packets from the input interface 100. On arrival of a packet, the queue memory manager
25 allocates a unique identifier to the packet, and sends the identifier for the packet to the stream identifier 102 on line 140. The queue memory manager additionally stores the packet in a temporary packet buffer. The queue memory manager also assigns a reference counter to the packet, and initialises the reference counter for that packet to zero. As will be described further hereinbelow, the reference
30 counter is used by the queue memory manager to determine whether the packet associated with the reference counter is still being processed by the multiplexer 100, or whether it should be removed from (or not initially entered into) the queue memory 112.

35 The packet identifier is an identifier which uniquely identifies a packet. The packet identifier uniquely identifies the packet for the purpose of storing the packet

5 identifier in the queue memory 112 as is described further hereinbelow and distinguishing each packet from other packets in the memory. In a simple implementation the packet identifier is a number, each packet being allocated a packet identifier being the next number in the sequence. Alternatively, the packet identifier may be composed of unique information from the packet header. In
10 systems in which the packet size is variable, the length of the packet may be included in the packet identifier.

The stream identifier 102 also receives the packets from the input interface on lines 118, and receives the packet identifier for each packet on line 140 from the
15 queue memory manager 114. The stream identifier 102 is responsible for determining which stream each packet belongs to. Thus, for example, the stream identifier will determine whether a particular packet is associated with a voice over IP stream, a block data transfer stream, or a telnet traffic stream. In accordance with the stream to which the particular received packet identifier belongs, the
20 stream identifier forwards the packet identifier for that packet to one of the policer/shapers 104 for further processing.

As will become apparent from the following description, the remainder of the processing is based on the packet identifier and not the packet. The packet
25 identifier advantageously provides an efficient representation of the packet. As will be described hereinafter, the queue memory and packet identifiers ensure that the original position in sequence of each individual packet is not lost in the multiplexing operation.

30 Policers/shapers 104 are an operational variant of a FIFO queue, in which there is subsidiary processing associated with the insertion and removal of elements from the queue. Such FIFO queues are well-known, and their implementation will be well within the scope of a person skilled in the art. The configuration of the policer/shapers is implementation dependent.

5 In the simple example shown in Figure 1 it is assumed that one policer/shaper is allocated to each of the different types of traffic stream being received. Thus, for example, the policer/shaper 104a is allocated to voice over IP traffic, the policer/shaper 104b is allocated to block data transfer traffic, and the policer/shaper 104c is allocated to telnet traffic. The policer/shaper 104d is shown
10 in Figure 1 by way of illustrating a means for managing packets associated with streams other than those processed by the policer/shapers 104a to 104c. Thus the stream identifier 102 forwards packet identifiers associated with voice over IP packets on line 120a to the policer/shaper 104a, forwards packet identifiers associated with block data transfer packets on line 120b to the policer/shaper
15 104b, and forwards packet identifiers associated with telnet traffic packets on line 120c to the policer/shaper 104c.

In the case of multi-casting or other services requiring replication (e.g. monitoring), the stream identifier 102 may forward packet identifiers to more than one
20 policer/shaper 104.

It should be noted that the stream identifier 102 does not forward packets for further processing within the multiplexer 100. Rather, the stream identifier forwards the packet identifiers allocated to packets for further processing in the
25 multiplexer 100.

The policer/shapers 104 are responsible for the assignment of quality classifications to the packets within a stream, and responsible for the quality control of the stream in both the short term and the long term by selectively
30 discarding packet identifiers in the stream and selectively time-shifting packets in the stream.

One function of policer/shapers is to service packets, using the corresponding packet identifiers, with a variable service rate. As a result, the packet identifiers
35 leaving a policer/shaper are variably spaced, preferably with a random or pseudo-random spacing. Spacing the packets randomly ensures that the least urgent

5 traffic is eventually serviced by the cherish/urgency multiplexor 106. The variable spacing of the packets at the output of the policer/shapers 104 reduces the coherence between streams from independent sources. Creating independent temporal patterns between streams increase the fairness of the cherish/urgency multiplexor decision processes. The cherish/urgency classifications allocated to
10 the packet identifiers in the policer/shapers do not on their own ensure fairness unconditionally.

For example, the delay experienced by a packet with an associated low urgency level depends on the temporal pattern of the more urgent traffic streams. If two
15 streams with different cherish/urgency levels are temporarily coherent, the more cherished packets may always "win" the race to enter the cherish/urgency multiplexor 106 and/or the more urgent data may always be transmitted first. This becomes statistically less likely if the streams have variably spaced packets.

20 Referring to Figure 4, there is illustrated two of the policer/shapers 104a and 104b of Figure 1, with examples of data packet streams on their inputs and outputs. The policer/shaper 104a receives packets 200 at a constant spacing t on its input line 120a. The policer/shaper 104b receives packet identifiers 202 at a constant spacing t on its input line 120b. The policer/shaper 104a services the packet
25 identifiers in its queue such that the packets 200a, 200b and 200c are generated on its output line 122a with variable spacing therebetween. Similarly, the servicing of packet identifiers by the policer/shaper 104b results in packets identifiers 202a, 202b and 202c on its output line 122b having variable spacing.

30 A second function of the policer/shapers is to limit the volume of traffic on a stream, achieved by the selective discard of packets. The nature of the policing policy is specified in terms of a load profile, which is a function of (at least) the offered load (i.e. the rate of arrival of packets on a stream). Therefore, and advantageously, there can be different characteristics defined for various load
35 levels. That is, quality can be down-graded or up-graded as the arrival rate increases. The transport characteristics can therefore be configured to better

- 5 match the application requirements and to avoid the effects of abnormal offered load, e.g. denial of service attacks. Referring to Figure 7, the principle of defining transport characteristics in terms of offered load is illustrated.

- 10 In Figure 7 the X-axis illustrates the offered load and the Y-axis illustrates the load which is actually transported. As can be seen in Figure 7 there are two threshold levels: a lower threshold level 250; and an upper threshold level 252. If the offered load does not exceed the threshold level 252, then the offered load can be transported within configured guarantees. If the offered load exceeds the threshold level 252, then the offered load cannot be guaranteed to be transmitted.
- 15 The lower threshold level 250 represents the minimum transported load regardless of how high the offered load becomes.

- Thus referring once again to Figure 7, once the offered load reaches the upper threshold level 252 and point 254, the transported load is adjusted such that it is reduced. As the offered load increases, the transported load continues to reduce to the point where it steadies out and tends towards the lower threshold level 250.
- 20

- Thus, the rate of packet discard being determined by the offered load. The upper threshold level defines a level at which the selective discard of packets is triggered. Packets are preferably discarded based on an instantaneous approximation of the offered load. Packets may be discarded probabilistically based on an instantaneous approximation of the offered load. If the offered load results in a transported load exceeding the upper threshold level, the transported load is reduced below the upper threshold level by selectively discarding packets.
- 25
- 30 When reduced, the transported load is preferably reduced to a level above the lower threshold level. The reduction in the transported load is greater the larger the offered load. The transported load is further reduced responsive to an increase in the offered load.

- 35 A Policer/Shaper achieves a desired load profile by selective discard of packets. Packets are preferably discarded based on an instantaneous approximation of the

5 offered load. Referring to Figure 10, there is illustrated a load profile for an exemplary implementation of a Policier/Shaper. Figure 10 illustrates a load profile for an exemplary policer/shaper implemented as a queue with 10 buffers and stochastic service times sampled from an exponential distribution, with rate parameters $\mu_1=\mu_2=\dots=\mu_9=1.3234$, $\mu_{10}=0.5$. The graph exhibits a load profile for a
10 Policier/Shaper implemented as a queue with 10 buffers having service times sampled from an exponential distribution whose rate parameter is dependent on the number of packets queuing. A description of such an implementation is provided in an example hereinafter. The service rate parameters are $\mu_1=\mu_2=\dots=\mu_9=1.3234$, $\mu_{10}=0.5$, the numeric values for the service rates and
15 offered and transported load being scaled so that the upper threshold level takes the value 1. The load profile illustrated in Figure 10 depicts the average transported load when the arriving traffic has a Poisson pattern. The application of standard queuing theory results permit the determination of the expected loss rate.

20

A third function of the policer/shapers 104 is to allocate cherish and urgency classifications to the packet identifiers included in their queues. By allocating such classifications to the packet identifiers, the packets themselves are inherently allocated the same classification. The principle of cherish and urgency
25 levels is discussed in international patent application no. PCT/GB00/01569 (PWF Ref. 101926). A cherish level indicates a class or level of loss for a packet. The class of loss indicates the tendency of the packet to be discarded. An urgency level indicates the level of urgency by which a packet should be processed. The urgency and cherish levels can, in combination be considered to constitute a
30 priority level for a packet. Thus a priority level has two components in this context.

As discussed hereinabove, in this exemplary policer/shaper, transport characteristics for a stream, including allocation of cherish and urgency levels, are
35 determined based on the offered load. The cherish and urgency classifications are advantageously assigned simultaneously to packet identifiers based on a

5 function of the current state of the queue when the classification is being
calculated. The likelihood of being in a particular state is a function of the offered
load. This classification function is configurable, and can be chosen without
constraint. For example, a given policer/shaper may be configured to assign one
of two classifications to packets based on a probabilistic choice. The classification
10 probability used may be related to the length of the queue. That is, such a
configuration may be designed to allocate higher classifications to packets with a
higher probability when the offered load is low.

The classification of a packet determines the maximum loss and delay it will most
15 likely experience when it is multiplexed with other streams at the output of the
policer/shapers, as will be described further hereinbelow. This is a separate
concept to that of the loss and delay the packet may experience inside the
policer/shaper itself. The loss experienced by a packet in the policer/shaper
depends on the recent arrival rate of the stream and the length of the queue. The
20 delay is determined by the configured service rates and the length of the queue.

Each of the policer/shapers 104 is preferably embodied using a queue with a
variable service rate. When a policer/shaper 104 receives a packet identifier, it
determines whether the packet identifier should be stored in its internal queue or
25 discarded. The control of the admission of a packet identifier into a policer/shaper
queue is discussed further hereinbelow.

In the following, the operation by which packets are admitted to or discarded from
policer/shapers 104 is first discussed. For the purposes of this discussion, an
30 example is taken of a policer/shaper utilising a queuing system having a queue
length of four as illustrated in Figure 2. Each state in Figure 2 is labelled with the
length of the queue in that state. The service times in each state are obtained by
sampling an exponential distribution with rate parameter μ . The service times
determine the conceptual rate of service for the packet. Two different service
35 rates are used in the example shown in Figure 2. If the queue is in state 1 (ie the
queue has only one packet) then μ_1 determines the rate of service for the packet

5 at the head of the queue. If the queue is longer, for example in state 3, service rate μ_2 would be used. This example embodiment of the policer/shaper also has arbitrary discard probabilities associated with each state. That is, on arrival, there is a probability that the packet identifier will be arbitrarily discarded. In this example, the probability of this event depends on the state of the queue when the
10 packet identifier arrives. The operation of a policer/shaper is now described in terms of the state of the queue on arrival of the packet identifier.

In a first case, it is assumed that the queue is full at the time the packet identifier arrives. In such a case, the probability of the packet identifier being entered into
15 the queue is zero. This does not necessarily mean that the packet is automatically discarded, as the packet may have been sent to more than one of the available policer/shapers. In such a case, the policer/shaper 104 forwards a command on its respective line 138 to decrement the reference count for the packet associated with a particular identifier. The queue memory manager then will discard the
20 packet referenced by the packet identifier if its reference count is 0 or less.

In a second case, at the time of arrival of a packet identifier at the policer/shaper 104 the queue is nearly full. For example, suppose the queue is in state 3. In state 3, there is a 30% chance that the packet identifier will not be entered into the
25 queue. If the packet identifier is admitted, then it is stored in the queue of the policer/shaper and the queue moves to state 4. The queue is then full, and any packet identifiers arriving before a packet identifier in the queue departs will not be admitted.

30 If the packet identifier is entered into the queue of the policer/shaper 104, the policer/shaper 104 sends an appropriate signal on its respective line 138 to the queue memory manager indicating that the packet associated with the packet identifier is to be admitted into the queue memory 112 and stored in a specific queue allocated for packets belonging to this stream.

35

- 5 When a departure is scheduled, the packet identifier at the head of the queue is serviced. The sampled rate used to service this packet identifier depends on the state of the queue. On the basis that the queue is in state 4 the sample service rate used to service this packet identifier is determined by an exponentially distributed random variable with mean μ_2 . The calculated service time of this
10 packet is based on the sample service rate and the length of the packet.

In a third case a packet identifier arrives at a time when the queue is nearly empty. The processing of the packet identifier in this case is very similar to the case when the queue is nearly full. Suppose a packet identifier arrives when the
15 queue is in state 1. There is no chance of the packet being arbitrarily discarded, since the probability of this event is configured to be zero. Therefore the packet identifier is stored in the queue, and the queue moves to state 2. If a departure event is scheduled before another arrival occurs, the packet identifier at the head of the queue is serviced based on the service rate for state 2 which is μ_1 .

20

In a fourth case the packet identifier arrives at a time when the queue is empty. In this case, as in the third case, the packet identifier will be admitted to the queue.

The description contained hereinabove has assumed that the policer/shaper is not
25 configured to send any packets without delay. In other words, the policer/shaper preferably has no low load transparency, all packets being delayed.

In its simplest form, low load transparency allows the first packet arriving at an empty queue to be forwarded immediately. Notwithstanding this action, the queue
30 moves into state 1. Subsequent packets arriving are processed as though this packet is present, except that when this packet would normally have been forwarded on expiry of its timer, it is not sent. Whether or not a packet identifier has been forwarded immediately on arrival or not is recorded in the queue. This concept can similarly be extended to multiple packets.

35

- 5 Before a packet identifier leaves a policer/shaper, the policer/shaper determines a quality classification for that packet identifier. The policer/shapers classify the packet identifiers, as described hereinabove, with cherish and urgency classifications. Each packet identifier must be classified with a cherish and urgency classification before it is forwarded to the cherish/urgency multiplexer
- 10 106. The classification assigned to a packet identifier is a function of the current state of the queue.

Again assume a policer/shaper having a queue of length 4. As shown in Figure 2, it is assumed that the queue of the policer/shaper is configured such that the

15 service rates for states 1 and 2 are defined as μ_1 , and the arbitrary loss possibility for states 1 and 2 is zero. States 3 and 4 have a defined service rate of μ_2 , and have an arbitrary loss probability of 0.3 (ie 30%).

In a preferred embodiment, each policer/shaper is configured with a primary and a

20 secondary cherish/urgency classification and a packet identifier is assigned one of the classifications upon arrival. Each state has an associated probability of a packet being classified with the primary classification. The probability of classifying a packet with the primary classification in each state is, in this example, configured to be: for state 1, 100%; for state 2, 80%; for state 3, 60%; and for

25 state 4, 40%. These probabilities are illustrated in Figure 3.

For example, the primary classification may be a more desirable classification than the secondary classification. If the offered load on the stream is low, the packets have a higher probability of being assigned the more desirable primary

30 classification, as the queue will most often be in state 1 or 2. As the offered load increases, there is a higher probability that a packet identifier will be assigned the secondary classification, which is less desirable, as this could mean the packet identifier will experience more delay and/or loss in the cherish/urgency multiplexor.

- 5 Referring to the discussion hereinabove of the criteria for admitting a packet identifier to the queue, when a packet identifier is admitted to the queue in case 2 (ie in the example when the queue is in state 3 and the packet identifier is admitted, moving the queue to state 4) the probability of classifying the packet identifier with the primary classification is 40% as the queue is then in state 4.
- 10 Therefore there is a 40% chance that the packet identifier will be assigned the primary classification, and a 60% chance that it will be assigned the secondary classification. This classification of the packet identifier is based on a simple probabilistic choice.
- 15 A packet identifier is then emitted from the policer/shaper queue at the end of the calculated service time, and the queue moves to state 3. A new calculated service time based on the defined parameters of state 3 then determines when the next departure is performed. If an arrival occurs before this period of time expires, the queue may then move to state 4 again, based on whether the arrival
- 20 is arbitrarily discarded or not.

Referring to case 3 above, in the case where a packet identifier arrives when the queue is nearly empty, the packet identifier is classified using the probability associated with state 2. In this case, there is an 80% chance of the packet

25 identifier being sent to the cherish/urgency multiplexer 106 with the primary classification.

Arbitrarily discarding packets on their arrival at a policer/shaper not only reduces the arrival rate of the stream, but also helps to avoid burst loss. For example, if a

30 burst of six packets arrived in the queuing system of Figure 2, the last two packets would be lost if packets were not arbitrarily discarded. On the other hand, if the probability of arbitrary discard increases with the queue length, it may be the case that, for example, the fourth or third packet is discarded on arrival, thus distributing the loss in the burst more fairly.

5 Another function of the policer/shaper as described hereinabove, is to keep the packet associated with the packet identifier being processed in the queue memory 112 by interacting with the queue memory manager 114. Storing packets in queues is necessary to ensure that packets in a stream are not re-ordered by the service process of the cherish/urgency multiplexer 106. Depending on the
10 function chosen to assign cherish/urgency levels to packets, there is the possibility the packets could be re-ordered during the multiplexing onto line 124, as discussed further hereinafter.

As an example, consider the simple probabilistic classification function described
15 above. Assume a burst of four packet identifiers arrive in the policer/shaper and all packets are stored in the queue. Furthermore, assume no more arrivals occur during the first period in question. Furthermore again, assume that the cherish/urgency multiplexor is not empty, with packets originating from other policer/shapers. Finally, assume that the primary cherish/urgency classification
20 has a desirable high urgency level, while the secondary classification has a low urgency level. Given these conditions, the packet identifier at the start of the burst, referred to as "packet identifier 1" has a higher probability of being assigned the secondary cherish/urgency classification than a packet near to the end of the burst, called "packet identifier 4". If packet identifier 1 is assigned the secondary
25 classification, while packet identifier 4 is assigned the primary classification, the difference in urgency levels may cause packet identifier 4 to be serviced before packet identifier 1 in the cherish/urgency multiplexer 106.

In order to avoid this, policer/shapers instruct the queue memory manager 114 to
30 queue packets in the queue memory 112 according to the order in which the corresponding packet identifiers are received in the policer/shaper. That is, on arrival in the policer/shaper, if a packet identifier is not discarded the queue memory is instructed to queue the packet in the queue of the relevant stream in the order in which it arrived.

35

- 5 Figure 8 shows an example implementation of a policer/shaper 104. The exemplary Policer/Shaper includes Policer/Shaper arrival process block 701, a packet identifier queue block 702, a Policer/Shaper departure process block 704, a timer 703, a discard probability generator 706, and a service time generator 707.
- 10 Packet identifiers arrive via line 708 to the Policer/Shaper arrival process block 701. The Policer/Shaper arrival process block 701 notifies the Policer/Shaper configuration management block 705 via line 720 that a packet identifier has arrived. The Policer/Shaper configuration management block polls the packet identifier queue block 702 via line 724 to obtain the current length of the queue
- 15 within the Policer/Shaper. Based on the response from the packet identifier queue block via line 726, the Policer/Shaper configuration management block determines if the queue is full or not. If there is available capacity in the queue, the Policer/Shaper configuration management block 705 then determines whether or not the packet identifier should be arbitrarily discarded, using input from the
- 20 discard probability generator 706 on line 732.

If the packet identifier is to be admitted to the queue, the Policer/Shaper configuration management block 705 alerts the Policer/Shaper arrival process block 701 via line 722 to admit the packet identifier. On receipt of this response,

25 the Policer/Shaper arrival process block 701 sends a request via line 736 (equivalent to one of lines 138) to the queue memory manager 114 to enqueue the packet in the one queue in the queue memory which is allocated for this stream. The Policer/Shaper arrival process then forwards the packet identifier to the packet identifier queue block 702 via line 710. The Policer/Shaper

30 configuration management block 705 calculates a new service time, based on input from the service time generator 707 on line 734 and the length of the packet, and sends this service time to the Policer/Shaper departure process block 704 via line 728. The Policer/Shaper departure process forwards the new service time to the timer 703 via line 716. The timer 703 resets itself to wake up at the end of the

35 new service time.

5 If the Policer/Shaper configuration management block 705 determines that the queue is full, it instructs the Policer/Shaper arrival process block 701 via line 722 to discard the packet identifier. In this case, the Policer/Shaper arrival process block 701 sends a discard instruction to the queue memory manager 114 via line 736. The Policer/Shaper arrival process block then discards the packet identifier.

10

When the timer 703 wakes up, it sends a request via line 718 to the Policer/Shaper departure process block 704 to emit a packet identifier. The Policer/Shaper departure process block sends a request to the Policer/Shaper configuration management block 705 via line 730 for a classification and a new service time. The configuration management block 705 polls the packet identifier queue block 702 via lines 724 and 726 to obtain the queue's current length. The configuration management block uses the current length of the queue to determine the classification for the packet identifier which is about to be emitted. The classification is sent to the Policer/Shaper departure process block via line 728. The Policer/Shaper departure process 704 concatenates a queue identifier, specifying which queue in the queue memory 112 is used for storing packets of this stream, and the classification to the packet identifier and forwards this tuple of data on line 738.

It should be noted that in this example implementation the classification of the packet identifiers is carried out as the packets leave the queue. The point at which the packets are classified within the Policer/Shapers is, however, implementation dependent, and is not limited to the example given herein. As described earlier hereinabove, the packet identifiers may be classified on arrival rather than on departure.

30

If the queue identified in the packet identifier queue block 702 is non-empty, the Policer/Shaper configuration management block 705 also sends a new service time to the Policer/Shaper departure process block 704 via line 728. This service time is forwarded to the timer 703 via line 716, and the timer sets itself to wake up after this time. If the queue is empty, no action is taken.

35

- 5 Further possible modifications to the policer/shapers 104 shown in Figure 1 are illustrated with reference to Figure 6. As shown in Figure 6, the policer/shapers may be cascaded in various ways.

Referring to Figure 6(a), there is illustrated an advantageous arrangement in which the output of a policer/shaper 210 provides an input to two parallel policer/shapers 212 and 214. The policer/shaper 210 receives packet identifiers on line 222 from the stream identifier 102. In accordance with the standard operation of the policer/shaper, the policer/shaper 210 selectively discards packet identifiers and time-shifts packet identifiers to provide a modified flow of packet identifiers on its output. In a first embodiment the thus modified flow of packet identifiers is broadcast on lines 234 and 236 to each of the parallel policer/shapers 212 and 214. The respective policer/shapers 212 and 214 then selectively discard packet identifiers and time-shift packet identifiers to generate two further modified packet identifier flows on their outputs 224 and 226. In a second embodiment, the output of the policer/shaper 210 is selectively provided to either one or the other of the policer/shapers 212 and 214. The appropriate one of the policer/shapers 212 and 214 then selectively discards packet identifiers and time-shifts packet identifiers onto its respective output.

- 25 Referring to Figure 6(b), there is illustrated a further arrangement in which two parallel policer/shapers 216 and 218 receive inputs comprising packet identifiers on flows on lines 228 and 230. The output of both policer/shapers 216 and 218 on their respective outputs 238 and 240 form an input to a further policer/shaper 220. The policer/shaper 220 outputs packet identifiers on line 232.

30

- It will be apparent to the person skilled in the art how policer/shapers can be cascaded in various combinations of the arrangements shown in Figure 6. For example, the arrangements of Figures 6(a) and 6(b) may be cascaded. In addition, any one of the policer/shapers may also receive inputs from other sources. For example, the policer/shaper 212 of Figure 6(a) may receive an additional input which is not derived from another policer/shaper. One skilled in
- 35

5 the art will appreciate how various cascading arrangements may be implemented. The only constraint is that the cascaded policer/shapers must be connected in an acyclic graph.

10 On departure, the policer/shaper sends both the packet identifier, and its associated queue identifier and classification to the cherish/urgency multiplexor 106 via a respective output line 122. The cherish/urgency multiplexor requires the packet identifier in addition to the queue identifier in case it has to issue a discard instruction to the queue memory manager 114, as will be discussed in further detail hereinbelow.

15 The cherish/urgency multiplexor 106 manages the contention for the network resources between two or more streams. The cherish/urgency multiplexor 106 receives packet identifiers from the various policer/shapers 104, each packet identifier being tagged with a cherish/urgency classification.

20 The classification of a packet identifier defines a cherish level and an urgency level for the packet with which it is associated. The cherish/urgency multiplexor manages the contention between two or more streams by servicing packets (via their packet identifiers) depending on their urgency level and, when necessary, discarding packets depending on their cherish level. The cherish level determines which packets are entered into the cherish/urgency multiplexor. The urgency level determines the order in which the packets are taken from the cherish/urgency multiplexor. Classifying packets using cherish and urgency levels is described in International Patent Application No. PCT/GB00/01569 (PWF ref 101926). Thus
25 the cherish/urgency multiplexor 106 manages the contention between the three streams 122a, 122b and 122c at the outputs of the policer/shapers 104a, 104b and 104c.

30 When a packet identifier with its associated cherish/urgency classification arrives, the cherish/urgency multiplexor 106 determines whether the packet identifier should be stored or discarded. This is determined by the available storage

- 5 capacity in the cherish/urgency multiplexor and the cherish level associated with the identifier.

In one known method of implementing cherishing, access to a buffer in the cherish/urgency multiplexor is determined by associating a cherish level with a particular area of the buffer. That is, suppose a packet can be classified with any one of N cherish levels. Then, assume there are K_N buffers in total. Traffic of cherish level 1 can enter any of the K_N buffers, while traffic of cherish level i can only enter the first K_{N+i+1} buffers, where $K_0 = 0 < K_1 < \dots < K_N$. Therefore, packets of cherish level 1 will have greater probability of finding available buffer resources than packets of cherish level i , where $i > 1$.

A scheme is also necessary to determine which packet identifier to forward, that is, how the urgency levels should be used to determine the order in which packets are serviced. One approach is to forward the identifier at the head of the queue with the highest urgency level, although other approaches are possible.

The loss rate for packet identifiers in a given stream is determined both by the cherish level assigned to the packets and by the maximum arrival rate for all other streams which have a cherish level of equal or greater value. The order in which packets are serviced in the cherish/urgency multiplexor depends on their urgency levels and the arrival rates and patterns of more urgent traffic. This determines the delay experienced by the packet identifier in the cherish/urgency multiplexor.

If the cherish/urgency multiplexor discards a packet identifier it instructs the queue memory manager via line 136 to decrement the reference count for the appropriate packet from the queue memory 112.

The cherish/urgency multiplexor records statistics about the traffic flows, including but not limited to, the number of the bytes and/or the number of packets both accepted and rejected for each urgency level and each cherish level.

5 The cherish/urgency multiplexor forwards one of the stored packet identifiers along with its associated queue identifier from its internal buffers to the rate limiter 108 on line 124 responsive to a request from the rate limiter on line 126. The choice of which identifier pair is forwarded to the rate limiter responsive to a request therefrom is determined by the packet identifier's urgency level and the
10 selection mechanism of the cherish/urgency multiplexor.

An example implementation of a cherish/urgency multiplexor 106 is illustrated in Figure 9. The cherish/urgency multiplexor includes a cherish/urgency arrival process block 801, a set of urgency queues 814 comprising queues 802 to 805, a
15 cherish/urgency departure process block 816, and a cherish/urgency configuration management block 807.

Packet identifiers arrive via line 810 to the Cherish/Urgency arrival process block 801. The Cherish/Urgency arrival process block notifies the Cherish/Urgency
20 configuration management block 807 that a packet identifier has arrived and forwards its cherish level via line 822. The Cherish/Urgency configuration management block 807 requests the length of each one of the urgency queues 802-805 via line 826. Based on the current total number of packet identifiers queueing in the cherish/urgency multiplexor 106 and the cherish level of the
25 packet identifier, the Cherish/Urgency configuration management block determines whether or not to discard the packet.

If the packet identifier is to be discarded, the Cherish/Urgency configuration management block 807 notifies the Cherish/Urgency arrival process 801
30 accordingly via line 824. In this case, the Cherish/Urgency arrival process sends an instruction to discard the packet to the queue memory via line 812, identifying the packet to discard by its packet identifier and its queue identifier. These identifiers are then discarded by the Chersih/Urgency arrival process block 801.

35 Otherwise, the Cherish/Urgency configuration management block 807 notifies the Chersih/Urgency arrival process 801 via line 722 to forward the packet identifier to

5 one of the urgency queues 802-805. There is one urgency queue for each
urgency level. This example implementation illustrates a cherish/urgency
multiplexor configured for four possible urgency levels. The Cherish/Urgency
arrival process obtains the urgency level for the packet identifier and forwards the
packet identifier and its associated queue identifier via one of the lines 830 to the
10 appropriate urgency queue.

Requests for packets from the rate limiter are received by the Chersih/Urgency
departure process block 806 via line 818. When a packet request arrives, the
Chersih/Urgency departure process requests the head element of one of the
15 urgency queues 802-805 via one of the lines 834. In the preferred embodiment,
the identifiers at the head of the most urgent queue will be requested. The packet
identifier and queue identifier pair at the head of the queue which has received the
request are forwarded to the Chersih/Urgency departure process block 806 via
one of the lines 832. The Chersih/Urgency departure process block 806 forwards
20 the pair of identifiers immediately to the rate limiter via line 124.

The rate limiter 108 moves contention from a point downstream in the network to
within the multiplexor 100 by restricting the service rate to one for which the
network has sufficient resources. The rate limiter ensures that the maximum
25 service rate is not exceeded in the long term by assuring that appropriate gaps
are maintained between transmitted packets.

The rate limiter 108 requests packet identifiers from the cherish/urgency
multiplexor on a request line 126, and receives packet identifiers along with their
30 associated queue identifiers on line 124. On receiving the pair of identifiers on
line 124, the rate limiter provides the queue identifier on line 134 to the queue
memory manager 114. Responsive thereto, the queue memory manager 114
provides the packet at the head of the identified queue from the queue memory to
the rate limiter on line 132. The rate limiter then forwards packets for transmission
35 to the output interface on line 128. On forwarding a packet to the output interface,
the rate limiter sets a timer to represent the service time of the particular packet at

5 the configured rate. At the end of the timing period assigned to the servicing of the particular packet, the rate limiter requests a further packet on line 126 from the cherish/urgency multiplexor 106.

10 If no re-ordering has occurred in the cherish/urgency multiplexor, the packet sent from the queue memory manager will be the same packet identified by the packet identifier. Otherwise, the packet identifier received by the rate limiter will refer to a packet which is still waiting in the queue in the queue memory.

15 The rate limiter can service packets stochastically or deterministically. The rate limiter may thus service packets such that they have variable spacing, as discussed hereinabove with reference to the policer/shapers.

20 A unit can be made from combining a cherish/urgency multiplexor 106 and rate limiter 108, preferably one which services packets stochastically. Such units can be cascaded with a plurality of policer/shapers and other such cherish/urgency multiplexor and rate limiter units and may receive additional inputs from other sources. One skilled in the art will appreciate how various cascading arrangements may be implemented. The only constraint is that the cascaded combination of these units must be connected in an acyclic graph.

25

If stochastic service rates are used, there is the potential that the sampled service rate may be much faster than the rate at which the packet can be physically transmitted by the output interface. In this case, the calculated resource size of the packet will be smaller than the actual size of the packet as illustrated in Figure 5. Referring to Figure 5a, there is illustrated the actual size of the packet, according to the physical transmission rate of the transmission medium to which it will be forwarded. Figure 5b illustrates the size of the packet where the service rate chosen for the packet is greater than the rate at which the packet will actually be transmitted on the physical medium. Figure 5c illustrates the size of the packet in the case where the service rate chosen for the packet is smaller than the rate at which the packet will actually be transmitted by the output interface 110.

30

35

5

As can be seen from Figure 5, if the calculated resource size is actually smaller than the actual resource size of the packet, two packets may overlap. In a practical system two packets cannot overlap during transmission. Instead, a burst of two or more packets will be observed. Therefore, when the packet identifiers are serviced stochastically, there is a finite probability that two or more packets will be sent back-to-back.

The output interface 110 provides flow control feedback on line 130 to the rate limiter 108. If the output interface indicates to the rate limiter that transmission has been suspended, the internal timing mechanism of the rate limiter is affected. In this case, the rate limiter can perform one of several actions. For example, the rate limiter may be configured to discard any packet which is scheduled for servicing during the suspension period. Such discarding of packets by the rate limiter requires the rate limiter to generate an alarm signal.

20

The rate limiter 108 can also be optimised to maximise the resource of the external system at a low load. That is, under low loads, the service rate can be chosen which is faster than the configured rate. As the load on the system increases, the service rate will be decreased until this is near to or at the configured rate.

25

The output interface 110 provides the basic functionality necessary for the onward transmission of data packets. The output interface 110 provides flow control feedback on lines 130 to the rate limiter 108 when some external back pressure (for example from the transmission medium) occurs. Thus the output interface transmits packets on line 142, and receives flow control signals on line 144 from the operating system or device driver.

30

The output interface is not responsible for any direct processing relating to the quality management of the data stream. If a packet is received by the output interface and the external system or device indicates that it has available

35

- 5 resources for transmission, then the packet is transmitted by the output interface without delay. If the rate limiter is operating in deterministic mode and flow control has not been exerted, there is never more than one packet in the buffer of the output interface at any one time.
- 10 Thus there has been described an invention which may be advantageously utilised in a multiplexor for guaranteeing a predetermined quality of service. Although the invention is described herein with reference to a particular exemplary embodiment, the person skilled in the art will appreciate that the invention is not so limited in its applicability, and may be more generally applied.

15

20

5 CLAIMS

1. A method of controlling an information flow in a data transmission system comprising: receiving a plurality of data packets; selectively discarding and/or selectively time-shifting the plurality of data packets in a predetermined manner to
10 thereby generate a modified output flow; selectively discarding and/or selectively time-shifting the data packets of the modified output flow in at least one of a plurality of further parallel steps to thereby generate at least one further modified output flow.
- 15 2. The method of claim 1 in which the modified output flow is further modified in all of the plurality of further parallel steps.
3. The method of claim 1 in which the modified output flow is further modified in a selected one of the plurality of further parallel steps.
20
4. The method of claim 1 in which the modified output flow is further modified in selected ones of the plurality of further parallel steps.
5. The method of any one of claims 1 to 4 in which the step of selectively
25 discarding and/or selectively time-shifting packets further comprises generating a modified output flow in which the packets are variably spaced.
6. The method of claim 5 in which the flow of packets is modulated to thereby generate the modified information flow.
30
7. The method of claim 5 or claim 6 in which the variable spacing is random or pseudo-random.
8. The method of any one of claims 1 to 7 wherein the packets are received at
35 a received sample rate, each packet having a received packet length, the method

5 further comprising determining a configured packet length for each packet based on a configured service rate for each packet.

9. The method of any one of claims 1 to 8 wherein the received packets in each selectively discarding and/or selectively time-shifting step are selectively
10 admitted to a queue.

10. The method of claim 9 in which the probability of a packet being admitted to the queue is dependent upon the length of the queue.

15 11. The method of claim 9 in which the probability of a packet already present in the queue being discarded to permit the admittance of the received packet is dependent on the length of the queue.

12. The method of any one of claims 9 to 11, in which a level of service is associated with each packet in the queue, the level of service being determined
20 by the length of the queue.

13. The method of claim 12 in which the level of service is controlled by controlling packet admissions to the queue.

25 14. The method of any one of claims 1 to 13 further comprising allocating a priority level to each packet.

15. The method of claim 14 in which packets are serviced in dependence on the allocated priority level.

30

16. The method of claim 14 to 15 wherein the priority level includes an indication of urgency associated with the packet and a level of loss which can be tolerated for the packet.

5 17. The method of any one of claims 14 to 16 when dependent on any one of
claims 9, 10, 11 or 13 wherein the priority level is allocated to a packet based on a
function of the state of the queue.

10 18. The method of any one of claims 8 to 17 wherein if the configured packet
length of a packet is such that a negative spacing is determined between two
packets, the two packets are adjacently multiplexed.

15 19. A controller for controlling an information flow in a data transmission
system, the controller comprising: first control means for selectively discarding
and/or selectively time-shifting a plurality of data packets in a predetermined
manner to thereby generate a modified information flow; and a plurality of parallel
second control means for selectively discarding and/or selectively time-shifting the
modified information flow in a predetermined manner, thereby to provide a further
modified information flow.

20

20. The controller of claim 19 in which at least one control means includes
modifying means for generating a modified information flow in which the packets
are variably spaced.

25 21. The controller of claim 20 in which the modifying means comprises a
modulator.

22. The controller of any one of claims 19 to 21 in which at least one control
means comprises a policer/shaper.

30

1 / 5




FIG. 1

2 / 5




FIG. 2

THE ARBITRARY LOSS PROBABILITIES AND THE SERVICE RATES
ASSIGNED TO THE STATES OF THE SHAPER / POLICER QUEUE




FIG. 3

THE CLASSIFICATION PROBABILITIES ASSOCIATED WITH
EACH STATE OF THE QUEUE




FIG. 4

3 / 5




FIG. 5





FIG. 6

4 / 5




FIG. 7




FIG. 8

5 / 5




FIG. 9




FIG. 10

INTERNATIONAL SEARCH REPORT

International Application No.

PCT/GB 00/03783

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/56 H04Q11/04

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04L H04Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>WU J -S ET AL: "GCRA-based architecture of multi-connection shaper and enforcer in multi-service ATM networks" COMPUTER COMMUNICATIONS,NL,ELSEVIER SCIENCE PUBLISHERS BV, AMSTERDAM, vol. 20, no. 8, 25 August 1997 (1997-08-25), pages 681-693, XP004126721 ISSN: 0140-3664 figures 4,8,10,11 page 682, column 1, paragraph 2.1 page 683, column 1, paragraph 3 page 685, column 2, paragraph 4.1.2 page 688, column 1, paragraph 5 -page 689, column 1</p> <p style="text-align: center;">— -/-</p>	1,3-5, 14-16,19

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

8 document member of the same patent family

Date of the actual completion of the international search

21 June 2001

Date of mailing of the international search report

28/06/2001

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax (+31-70) 340-3016

Authorized officer

Lamadie, S

INTERNATIONAL SEARCH REPORT

International Application No.

PCT/GB 00/03783

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 602 845 A (WAHL STEFAN) 11 February 1997 (1997-02-11) figures 1,2 column 5, line 23 - line 42 column 6, line 11 - line 28	5,20

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PT/GB 00/03783

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5602845 A	11-02-1997	DE 4442532 A CA 2163594 A EP 0715482 A	05-06-1996 31-05-1996 05-06-1996